



CLAIMS LISTING

Claims 1-41 cancelled.

42. A computer implemented method for integrating and transforming data, said method comprising:

creating a plurality of semantic models comprising at least a first semantic model and a second semantic model, each semantic model having a plurality of semantic components comprising:

a first component representing a first concept;

a second component representing a second concept; and,

a third component representing at least one relationship among the first concept represented by the first component and the second concept represented by the second component;

requiring the first semantic model's concepts and first semantic model's relationships to be those relevant to a first category of knowledge;

recording a first association among the first semantic model and the first category of knowledge;

requiring the second semantic model's concepts and second semantic model's relationships to be those relevant to a second category of knowledge;

recording a second association among the second semantic model and the second category of knowledge;

storing both the first semantic model and the second semantic model;

creating a first model mapping by automatically mapping at least the first concept of the first semantic model to at least the first concept of the second semantic model;

storing the first model mapping;

identifying a first source associated with the first semantic model;

identifying a first destination associated with the second semantic model;

accepting a first input from the first source and comprising a first input data;

determining that the first input data represents an instance of the first concept of the first semantic model;

determining that the first input data is to be transformed according to the first model mapping;

transforming the first input data consistent with the first model mapping to create a first output comprising a first output data;

determining that the first output data represents an instance of the first concept of the second semantic model;

specifying a set of validation rules providing a prescription for the valid use of data conforming to at least one of the first semantic model and the second semantic model;

validating at least one of the first input data and the first output data according to the set of validation rules;

writing, after the steps of transforming and validating, the first output data to the first destination; and,

communicating the first output data to a destination software system.

43. (structure of set of sm) A method as in Claim 1, wherein the first semantic model is a first domain semantic model, the second semantic model is a second domain semantic model, and further comprising the steps of:

determining that the first category of knowledge and the second category of knowledge have portions in common, and that each portion pertaining to a first application domain;

storing a third association between the first semantic model and the first application domain;

storing a fourth association between the second semantic model and the first application domain;

implementing the first domain semantic model incorporating a first representation of least a first domain concept pertaining to the first application domain;

implementing the second domain semantic model incorporating a second representation of least the first domain concept pertaining to the first application domain;

associating the first model mapping with the first application domain;

and,

implementing the first model mapping as a first mapping between at least the first domain concept as represented by the first domain semantic model and the first domain concept as represented by the second domain semantic model.

44. A method as in Claim 1, wherein at least one of the first semantic model and the second semantic model describes the semantics of at least one of a message between software applications, an Internet-based software service implemented using Web Service standards, a business document, an XML document, and a database.

45. A method as in Claim 1, the step of creating a first model mapping further comprising:

determining the first source;

determining that the first semantic model is associated with the source;

selecting and accessing said first semantic model;

determining a first mapping subject comprising at least one of a specific domain and a specific topic;

determining that the first mapping subject is associated with the source;

determining a destination;

selecting and accessing said second semantic model based on association with the destination;

displaying at least some portion of each of the first semantic model and the second semantic model;

accepting a third input both responsive and subsequent to said step of displaying, said third input selective of the first semantic model's first concept;

accepting a fourth input both responsive and subsequent to said step of displaying, said fourth input selective of the second semantic model's first concept;

establishing a semantic component association between the first semantic model's first concept and the second semantic model's first concept;

accepting input definitional of a first user-defined transformation rule consistent with the semantic component association;

accepting input definitional of a first user-defined validation rule comprising at least one of a rule for validating the consistency of source data with the first semantic model, a rule for validating the consistency of destination data with the second semantic model, and a rule for validating the use of the first user-defined transformation rule for transforming source data into destination data consistent with each of the first semantic model, the second semantic model, and the first model mapping;

editing of at least one member of a set comprising any semantic component, the semantic component association, the first user-defined transformation rule, and the first user-defined validation rule;

storing the first model mapping; and,

storing an association between the mapping subject and each of the first semantic model, the second semantic model, and the first model mapping.

46. A method as in Claim 1 further comprising a sub-method for semantic model augmentation comprising:

importing a set of semantic information into at least one of the first semantic model and the second semantic model, said set of semantic information comprising any of concepts, relationships among concepts, rules among data representing concepts, rules among concepts, and components; and,

applying the set of semantic information to at least one of the first semantic model and the second semantic model;

47. A method as in Claim 1 further comprising a sub-method for semantic model creation comprising:

creating a new semantic model comprising at least one of the first semantic model, the second semantic model, and a seed semantic model by importing a set of semantic information, said set of semantic information comprising any of (1) concepts, (2) relationships among concepts, (3) data types, (4) relationships among data types, (5) rules among data representing instances of concepts, and (6)

rules among any combination of some concept, some component, and some seed semantic model; and,

storing the new semantic model.

48. A method as in Claim 1 further comprising using a first Adapter to perform at least one member of a set of operations on a set of semantic information, said set of operations comprising importing, exporting, modifying, storing, accepting, communicating, reading, recording, and writing.

49. A method as in Claim 1, further comprising:
accepting a second input;

using the second input to determine at least partially a first modification of a modifiable semantic model;

identifying the modifiable semantic model as at least one of the first semantic model, the second semantic model, a referent semantic model, a semantic model template, and a seed semantic model; and,

modifying the modifiable semantic model consistent with the first modification.

50. A method as in Claim 1, wherein some semantic model is created from a pre-defined seed semantic model.

51. A method as in Claim 1 further comprising recording at least one transformation rule and using said transformation rule in conjunction with at least one of the first semantic model and the second semantic model.

52. A method as in Claim 1, wherein at least one of the first semantic model and the second semantic model is implemented as at least one of an ontology, a semantic network, a frame, and a knowledge base.

53. A method as in Claim 1, wherein at least one of the first semantic model and the second semantic model is represented in a standardized language for at least one of knowledge description and knowledge querying, the definition of said standardized language being both publicly accessible and approved by a plurality of businesses.

54. A method as in Claim 1, wherein, prior to the step of writing, the first input data is processed according to at least a first rule in order to accomplish at least one of the operations of data profiling, semantic mapping, semantic resolution, document abstraction, data cleansing, data normalization, data transformation, data validation, and semantic model validation.

55. A method as in Claim 1 further comprising executing a rule to automatically correct for failure of at least one member of a set comprising at least one of some validation rule and some transformation rule.

56. A method as in Claim 1 further comprising generating from an initial representation of any of a data schema, a business pattern, semantic model, a rule, a model mapping, and an association, a new representation that is more optimal for efficient runtime processing than the initial representation.

57. A method as in Claim 1 wherein the knowledge represented by some semantic model is independent of knowledge of one of any referent, any domain, and any topic; and that semantic model is used in conjunction with multiple referents if independent of

knowledge of any referent, with multiple domains independent of knowledge of any domain, and with multiple topics if independent of knowledge of any referent.

58. A method as in Claim 1 wherein the first input data is known to be valid and is used to perform at least one modification operation of a set of modification operations comprising modifying a semantic model template, a semantic model, a rule, a model mapping.

59. A method as in Claim 1 further comprising determining which of a plurality of semantic models applies to the first input by an analysis of at least one of the first input's content, name, format, structure, and type.

60. A method as in Claim 1 further comprising the step of automatically selecting, after the step of accepting the first input, the model mapping from a plurality of model mappings based on an analysis of any of the first input, the source, the destination, the first semantic model, the second semantic model, a referent semantic model, a semantic model template, a domain, and a topic.

61. A method as in Claim 1 wherein at least one rule is used to represent any of a semantic model, a model mapping, a data schema, a business pattern, a transformation, and a validation.

62. A method as in Claim 1 further comprising:
transforming the first data via a complex transformation; and,
using at least one of the inferencing methods of forward chaining and backward chaining to derive any of (1) the complex transformation and (2) an explanation of the complex transformation.

63. A method as in Claim 1 further comprising:
- automatically converting any of a data schema, a business pattern, a data profiling operation, a data normalization operation, a data cleansing operation, a validation, a transformation, a semantic model template, a semantic model, and a model mapping into a set of rules; and,
- applying at least one member of the set of rules to one of the first input and the first source.
64. A method as in Claim 1, further comprising creating at least a third semantic model and restricting the third semantic model to some category of knowledge.
65. A method as in Claim 64, wherein the third semantic model is a referent semantic model representing knowledge about at least one member of a referent set comprising at least one of a second source and a second destination.
66. A method as in Claim 1, further comprising a sub-method for interactive help, said sub-method for interactive help further comprising:
- aiding at least one operation of a set of operations comprising creating a semantic model, creating a model mapping, mapping one semantic model to another semantic model, creating a data validation rule, creating a semantic model validation rule, creating a data transformation rule, semantic model validation, data validation, data transformation, data normalization, data profiling, and data cleansing; and,
- responding to an event further comprising:
- identifying the event as a request for interactive help; and,

partially determining a response to the request from at least one of an identification of the one operation, the one operation's context, and the event's context.

67. A method as in Claim 66, wherein the sub-method for interactive help constitutes an Interactive Guide, said Interactive Guide comprising the step of:

aiding at least one member of a set of knowledge engineering operations comprising identifying a semantic concept within data, resolving semantic ambiguities, incorporating a semantic element within a semantic model, incorporating a relationship among a plurality of semantic elements within a semantic model, incorporating an association between an element of the first semantic model and an element of the second semantic model within the first model mapping, incorporating a validation rule, and incorporating a transformation rule.

68. A method as in Claim 66, wherein the sub-method for interactive help incorporates an Interactive Guide, said Interactive Guide comprising:

determining potential choices for incorporating at least one of a semantic element, a concept, a relationship, and an association;

computing said potential choices from at least one member of a set of heuristics, said set of heuristics comprising at least one of a syntactic heuristic, a structural heuristic, a human input heuristic, an a priori heuristic, and an inductive heuristic; and,

performing at least one member of a set of alterations comprising creating at least one new member of the set of heuristics, refining at least one member of the set of

heuristics, deleting at least one member of the set of heuristics, confirming at least one member of the set of heuristics, and denying at least one member of the set of heuristics.

69. A method as in Claim 1, further comprising:

detecting of at least one of an error, an inaccuracy, an invalidity, an inconsistency, and an inadequacy;

performing, responsive said step of detecting, at least one of modifying the model map, modifying the set of validation rules, modifying a set of transformation rules, and modifying some semantic model.

70. A method as in Claim 70 wherein the steps of detecting and performing, once initiated, are automatic and are completed without subsequent interactive input.

71. A method as in Claim 1, further comprising using a first Adapter to perform at least one member of a set of operations on a set of data comprising the first input data, said set of operations comprising importing, exporting, modifying, storing, accepting, communicating, reading, recording, and writing.

72. A method as in Claim 71, wherein the Adapter is a SOAP Message Handler compliant with a SOAP standard.

73. A method as in Claim 1, wherein the first semantic model is a first domain semantic model, the second semantic model is a second domain semantic model, and further comprising the steps of:

determining that the first category of knowledge and the second category of knowledge have portions in common, and that each pertaining to a first application domain;

identifying a set of topics pertaining to the first application domain and comprising at least a first topic;

storing a fifth association between the first topic and the first application domain;

implementing the first domain semantic model at least partially as a set of topic semantic models comprising at least a first topic semantic model; said first topic semantic model representing concepts pertaining to the first topic in the first application domain;

implementing the second domain semantic model at least partially as a set of topic semantic models comprising at least a second topic semantic model; said second topic semantic model representing concepts pertaining to the first topic in the first application domain;

associating the first model mapping with the first topic in the first application domain;

identifying the second topic semantic model within the second domain semantic model that represents the first topic; and,

implementing the first model mapping as a first mapping between the first topic semantic model and the second topic semantic model.

74. A method as in Claim 73, further comprising:
- identifying at least a first domain concept pertaining to a specific application domain;
- creating a set of template rules, said set of template rules comprising at least one rule, and said rule comprising any of a data validation rule, a semantic model validation rule, and a transformation rule;
- creating a first semantic model template for the specific application domain comprising identification of the specific application domain, the first concept, the first association; the set of topics, and the set of template rules;
- using the first semantic model template to represent knowledge of the specific application domain, wherein all knowledge represented by the first semantic model template is logically independent of any referent; and,
- storing the first semantic model template in a repository.
75. A method as in Claim 74, wherein the first semantic model template further comprises:
- a second topic distinct from the first topic; and,
- a set of relationships comprising at least a first relationship among the first topic semantic model's concepts and the second topic semantic model's concepts.
76. A method as in Claim 74, further comprising using the first semantic model template to perform at least one of creating, validating, and modifying some semantic model representing concepts in the first application domain.

77. A method as in Claim 74, wherein the first semantic model template further comprises at least one pre-defined rule applicable to at least one of the first input data represented by the first component of the first semantic model and the first output data represented by the second component of the second semantic model.
78. A method as in Claim 74 wherein the first semantic model template is used to validate one of the first input data, the first output data, the first semantic model, and the second semantic model.
79. A method as in Claim 74 wherein the knowledge represented by the first semantic model template is independent of knowledge of one of any referent, any domain, and any topic; and the first semantic model template is used in conjunction with multiple referents if independent of knowledge of any referent, with multiple domains independent of knowledge of any domain, and with multiple topics if independent of knowledge of any referent.
80. A method as in Claim 1, further comprising the steps of:
creating a candidate mapping set comprising at least one candidate mapping between elements of said first semantic model and said second semantic model;
assigning a weight to each candidate mapping in the candidate mapping set, said weight computed from a combination of one or more component weights, and each component weight is individually computable from some computable procedure;
establishing a validity criterion that is computationally dependent on the weight;

classifying each candidate mapping in the candidate mapping set, according to pre-defined validity criteria, as belonging to one of a set of validity classes comprising valid, invalid, and undetermined;

performing, subsequent to the step of classifying, any of a set of actions comprising (1) including in the candidate mapping set at least one candidate mapping that is valid and (2) excluding from the candidate mapping set at least one candidate mapping that is invalid;

displaying the candidate mapping set;

accepting selection of at least one candidate mapping in the candidate mapping set as the selected candidate mapping; and,

modifying the first model mapping according to the selected candidate mapping.

81. A method as in Claim 80, further comprising:

computing the weight;

determining that at least a first weight parameter on which the weight is dependent has changed and that the weight is dependent on the first weight parameter;

identifying a set of unchanging components weights, each member of which cannot change on recalculation in consequence of at least the first weight parameter being changed, and comprising at least a first component weight;

recomputing each component weight dependent on the first weight parameter; and,

recomputing the weight without recalculating any member of the set of unchanging component weights.

82. A method as in Claim 80, further comprising:
selecting a first selected candidate mapping;

generating output explanatory of how at least a portion of the weight of the first selected candidate mapping was computed.

83. A method as in Claim 80 further comprising at least one of (1) modifying at least one component weight and (2) modifying some computable procedure by which at least one component weight is computed.

84. A method as in Claim 80 further comprising:
identifying a first relationship that is any of descendant, sibling, child, parent, and ancestor;

automatically identifying a second element of the first semantic model having the first relationship to the first element of the first semantic model;

automatically identifying a second element of the second semantic model having the first relationship to the first element of the second semantic model;

automatically creating an additional candidate mapping between the second element of the first semantic model and the second element of the second semantic model;

automatically incorporating the additional candidate mapping in the candidate mapping set.

85. A method as in Claim 80, wherein the step of assigning a weight to each candidate mapping further comprises:

at least partially determining the weight according to a heuristic rule set, said heuristic rule set comprising at least a first heuristic rule; and,

using said first heuristic rule to determine at least a first component weight of the weight.

86. A method as in Claim 85, further comprising:

accepting input to perform at least one action in a set of actions comprising defining the first heuristic rule in the heuristic rule set and modifying the first heuristic rule in the heuristic rule set.

87. A method as in Claim 85, further comprising:

pre-defining the first heuristic rule; and,

determining, based on input, a criterion of applicability of the first heuristic rule to the step of determining potential choices.

88. A method as in Claim 80, wherein at least one candidate mapping is included in the candidate mapping set based on a function of at least one of (1) the weight of that candidate mapping and (2) a threshold value that is comparable to the weight of that candidate mapping.

89. A method as in Claim 88 further comprising modifying the threshold value.

90. A method as in Claim 80, wherein the number of candidate mappings included in the set is limited to a maximum number.

91. A method as in Claim 90, further comprising modifying the maximum number.

92. A general-purpose computer incorporating specific hardware and software in an architecture for transforming, profiling, cleansing, normalizing, validating, and integrating data in a flexible and continuously maintainable manner, wherein said specific hardware and software comprise:

- means for defining at least a first semantic model and a second semantic model;
- means for defining a model mapping among semantic models;
- means for storing said semantic models and said model mapping;
- means for defining validation rules and transformation rules;
- means for accepting data from at least one source;
- means for transforming said data according to the model mapping and transformation rules;
- means for validating said data according to the validation rules; and,
- means for moving said data to at least one destination.